

Demand Dispatch with Heterogeneous Intelligent Loads

Joel Mathias
Dept. ECE
Univ. of Florida

Ana Bušić
Inria Paris
Département d'Informatique de l'ENS

Sean Meyn
Dept. ECE
Univ. of Florida

Abstract A distributed control architecture is presented that is intended to make a collection of heterogeneous loads appear to the grid operator as a nearly perfect battery. Local control is based on randomized decision rules advocated in prior research, and extended in this paper to any load with a discrete number of power states. Additional linear filtering at the load ensures that the input-output dynamics of the aggregate has a nearly flat input-output response: the behavior of an ideal, multi-GW battery system.

Acknowledgements *Research supported by NSF grant CPS-1259040, and PGMO (Gaspard Monge Program for Optimization and operations research). We thank Rim Kaddah of Telecom Paristech for many stimulating conversations, especially relating to the design choices discussed in Sections 2.2 and 2.3*

1 Introduction

Are billion dollar batteries and billion dollar gas turbine generators required to manage the volatility of renewable generation?

In prior research, it is argued that balancing resources will come from flexible loads at much lower cost and potentially greater performance. In order to realize this vision, a decentralized control design is utilized; the design respects the limitations of the loads, which are based on dynamic constraints as well as strict bounds on the quality of service (QoS) delivered to consumers [6]. Automation is required to ensure that the grid operator obtains reliable ancillary service, and that reliability to the consumer is also maintained.

The goal of this research is to create virtual energy storage from flexible loads. The framework here and in prior research is *Demand Dispatch*: power consumption from loads varies in a possibly coordinated manner to automatically and continuously provide service to the grid, without impacting QoS to the consumer.

This paper investigates a question posed in [14]: *what intelligence is required at the grid-level to implement demand dispatch?* The question was addressed through extra layers of local control at each load. One topic left for future research was how to approximately invert the dynamics of an aggregate of loads so that the resulting dynamics would approximate a perfect battery.

The control architecture proposed in [14, 17] is illustrated in Fig. 1. The “compensation” block represents today’s balancing authority (BA), and the “grid” represents the aggregate dynamics of loads, generators, transmission lines, and other grid elements. Design of the compensator G_c will be based on an input-output model of the grid, denoted G_p in the figure [11] (see Section 2.1 for details). The contributions of this paper are summarized here:

- (i) New design techniques are introduced for a broad class of on/off loads (or more generally, loads with a finite number of operating states). It is found in numerical results that the resulting mean-

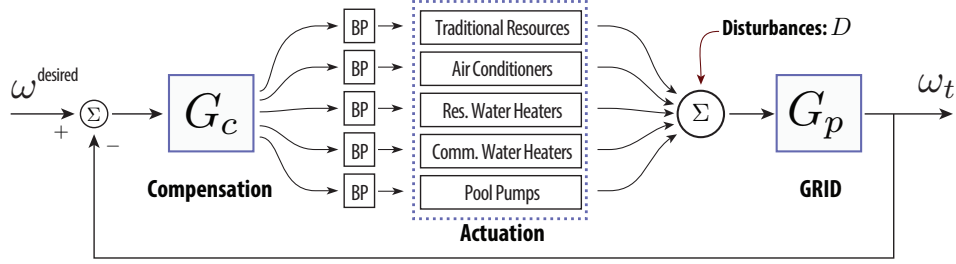


Figure 1: A control architecture for Demand Dispatch.

field dynamics share desirable properties observed in the optimal design of [17]. In particular, in every example considered, the linearized dynamics are minimum phase.

- (ii) It is argued that the minimum phase property is valuable in the design of a prefilter for each load. Applying techniques from the theory of robust control, it is shown that the input-output dynamics can be shaped to appear as a constant gain over a bandwidth centered at the nominal period of the load.
- (iii) With one-way communication from the BA to the loads, it is shown that the aggregate, with each load acting independently, serves as a nearly perfect “virtual battery”. This is argued based on control analysis, and tested through simulation.

Grid-level simulation experiments were conducted using over 50 different types of loads. Local control ensures that this diverse population — including pools with time constants of 24 hours, and residential air-conditioning with time constants of tens of minutes to one hour — act cooperatively to provide regulation over all time-scales.

Related research Beginning in the early eighties, deterministic schemes were introduced to model and control a population of thermostatically controlled loads (TCLs) for ancillary services [4, 13]. Randomized algorithms appeared in the sequels [10, 16]; system identification and state estimation are required for accurate tracking of the control signal.

Centralized control is the subject of [8], where the main contribution is to address combinatorial complexity through a priority ordering of loads. The state information required by the centralized controller presents challenges in terms of both communication and privacy.

Local control of refrigerators is proposed for primary frequency control in [21]. A randomized control architecture is introduced to avoid synchronization of loads. There are no performance guarantees with respect to ancillary service, which raises concerns as inaccurate primary frequency control can destabilize the grid [11, 14].

There is substantial literature on indirect load control, where customers are encouraged to shift their electricity usage in response to real-time prices. Control through price signals can introduce uncertain dynamics and present a risk to system stability [3, 18, 20].

Our approach simultaneously addresses the following four challenges: (i) a distributed control architecture simplifies communication infrastructure requirements and assuages consumer privacy concerns; (ii) local control ensures reliable ancillary service; (iii) local control also ensures QoS to the consumer; and (iv) contractual agreements and periodic credits, such as those proffered by Florida Power and Light in their OnCall program, are advocated to incentivize customer participation.

The present paper concerns all four challenges, but focuses on a new approach to topic (i). It is assumed that there is only one-way communication from BA to loads, and that the control signal generated by the grid operator is based on frequency deviation, akin to the manner in which AGC (automatic generation control) is synthesized today.




Fig. 1. ERCOT model elements.

Figure 2: Macro grid model of [5].

It is remarkable to see the potential for demand dispatch based on minimal communication. However, some communication from loads to the grid operator remains valuable in practice. In particular, the grid operator requires bounds on the capacity of service from loads, and may want to occasionally update or verify parameters in local control algorithms.

The remainder of this paper is organized as follows. Section 2 provides details of the distributed control architecture, including grid-level control, local control design, and load dynamics. In Section 3, multiple simulations are performed to demonstrate the validity and utility of demand dispatch. Conclusions and directions for future research are summarized in Section 4.

2 Distributed Control Architecture

2.1 Grid level control

The macro grid model used in this study is taken from [5], which is itself based on standard power systems analysis [11]. The grid is modeled as an input-output linear system whose input is power deviation and output, frequency deviation. The elements of their model are illustrated in Fig. 2, which is taken from Fig. 1 of [5]. A particular grid model of [5] is used in numerical experiments:

$$G_p(s) = 10^{-5} \frac{2.488s + 2.057}{s^2 + 0.3827s + 0.1071} \quad (1)$$

The impulse response of this system is in close agreement with the response of frequency to a grid outage in the ERCOT region — a full discussion can be found in [14].

Throughout the paper, the transfer function (1) is used to model the “grid block” shown in Fig. 1. The resonance of this transfer function corresponds to time-scales on the order of seconds, while in this paper the relevant disturbances to be rejected evolve on much slower time-scales. This justifies the use of PI control for the choice of G_c in the “compensation block”. The output of the compensator is denoted

$$U_t = K_P \tilde{\omega}_t + K_I \int_0^t \tilde{\omega}_r dr$$

where $\tilde{\omega}_t = \omega^{\text{desired}} - \omega_t$. The interpretation of U_t is the desired change in power from all resources (MWs). The control parameters (proportional and integral gains K_P , K_I) are chosen to respect the uncertainty of grid dynamics on timescales of seconds or faster.

In practice, the PI compensator would be modified to avoid “integrator windup”, as is common practice in synthesizing the AGC signal today.

For simplicity, in this paper, we focus solely on techniques for balancing within the control region (this includes ramp services, balancing reserves, and frequency regulation). In practice, the regulation of tie-line error would be performed in conjunction with these services.

The signal U is decomposed using several bandpass filters: Each block “BP” shown in Fig. 1 represents a bandpass filter that is chosen based on the dynamics and constraints of the associated aggregate of resources. Batteries and flywheels are valuable for the highest frequency component of

U ; demand dispatch based on refrigerators and water heaters can provide service on time-scales of tens of minutes to several hours [17].

2.2 Local control: Markovian dynamics

The local control described here is a continuous-time variant of the myopic design introduced in [2].

The starting point is the construction of a Markovian model for nominal behavior of an individual load. The state process evolves in continuous time, on a finite state space denoted \mathbf{X} . Hence its dynamics are defined by a *rate matrix*, denoted \mathcal{A}_0 . For two states $x, x' \in \mathbf{X}$ the transition probability is denoted $P^t(x, x') = \mathbb{P}\{X_t = x' \mid X_0 = x\}$, which is the matrix exponential $P^t = \exp(t\mathcal{A}_0)$, $t \geq 0$.

It is assumed that the nominal model has a unique invariant pmf (probability mass function), denoted π_0 . Invariance requires that $\sum_x \pi_0(x)\mathcal{A}_0(x, x') = 0$ for every $x' \in \mathbf{X}$.

The rate matrix is assumed to be of the following form,

$$\mathcal{A}_0 = r[-I + S_0] \quad (2)$$

where S_0 is a Markov transition matrix, and I is the identity matrix. A Markov process \mathbf{X} with rate matrix (2) can be realized by first constructing a Poisson process with rate r and jump times $\{T_k : k \geq 1\}$. The continuous-time process \mathbf{X} is constant on the inter-jump time-intervals $[T_k, T_{k+1})$, and

$$\mathbb{P}\{X_{T_{k+1}} = x' \mid X_{T_k} = x\} = S_0(x, x'),$$

for $x, x' \in \mathbf{X}$ and $k \geq 0$, with $T_0 = 0$. The assumption that $X_t = X_{T_k}$ for $t \in [T_k, T_{k+1})$ reflects the fact that we are only considering the load at the sampling times $\{T_k\}$. *We do not assume that the load state itself is constant over this period.*

The state of the load has the following form: $X_t = (X_t^u, X_t^n)$ for $t \geq 0$, and the state space has the form $\mathbf{X} = \mathbf{X}^u \times \mathbf{X}^n$. The first component \mathbf{X}^u represents a variable that can be adjusted directly, such as power consumption, or the temperature set-point for a refrigerator. The second component \mathbf{X}^n is *indirectly* controlled through \mathbf{X}^u and exogenous disturbances (e.g., someone opens the refrigerator).

The Markovian dynamics for the nominal model are assumed to be of the form

$$S_0(x, x') = R_0(x, x'_u)Q_0(x, x'_n), \quad (3)$$

where $x, x' \in \mathbf{X} = \mathbf{X}^u \times \mathbf{X}^n$, and $\sum_{x'_u} R_0(x, x'_u) = \sum_{x'_n} Q_0(x, x'_n) = 1$. The matrices R_0, Q_0 model the dynamics of $\mathbf{X}^u, \mathbf{X}^n$, respectively.

The construction of S_0 is of course entirely dependent on the characteristics of the particular load.

For simplicity, in this paper it is assumed that $X_t^u = m_t$ represents power consumption (that can be controlled directly at the load). It is assumed moreover that there are just two power states: on or off. The process \mathbf{m} evolves in the binary set denoted $\mathbf{X}^u = \{\ominus, \oplus\}$. Denote by $\mathcal{U}(x)$ the associated power consumption: $\mathcal{U}(\ominus, x_n) = 0$, and $\mathcal{U}(\oplus, x_n) = \varrho$ (kW) (a positive value, independent of $x_n \in \mathbf{X}^n$).

Local control is based on a perturbation of nominal behavior, defined by a family of rate matrices $\{\mathcal{A}_\zeta : \zeta \in \mathbb{R}\}$. The following myopic design is used in all of the numerical experiments considered here:

$$S_\zeta(x, x') = S_0(x, x') \exp(\zeta \mathcal{U}(x') - \Lambda_\zeta(x)) \quad (4)$$

in which $\Lambda_\zeta(x)$ is the normalizing constant defined so that $\sum_{x'} S_\zeta(x, x') = 1$ for each x .

The goal of the myopic design is to influence the load to consume more power at time t when $\zeta_t > 0$, and less power when $\zeta_t < 0$.

Given a homogeneous collection of N loads, the empirical distribution at time t is defined as follows:

$$\mu_t^N(x) := \frac{1}{N} \sum_{i=1}^N \mathbb{I}\{X_t^i = x\}, \quad x \in \mathbf{X}. \quad (5)$$

We assume this is approximated by the mean-field equations,

$$\frac{d}{dt}\mu_t = \mu_t \mathcal{A}_{\zeta_t} \quad (6)$$

in which μ_t is interpreted as a row vector; justification for large N is straightforward in the discrete-time setting [17]. The average power is denoted $y_t = \sum_x \mu_t(x) \mathcal{U}(x)$, and the steady-state average power consumption for the nominal model is $\bar{y}^0 = \sum_x \pi_0(x) \mathcal{U}(x)$.

It is assumed moreover that \mathcal{A}_ζ is continuously differentiable in ζ . This justifies the linear state space model approximation,

$$\frac{d}{dt}\Phi_t = A\Phi_t + B\zeta_t, \quad \gamma_t = C\Phi_t \quad (7)$$

where $A = \mathcal{A}'_0$, and B, C^T are column vectors of dimension $d = |\mathbf{X}|$:

$$C_k = \mathcal{U}(x^k), \quad B_k = \sum_x \pi_0(x) \mathcal{A}'_0(x, x^k), \quad 1 \leq k \leq d$$

where \mathcal{A}'_0 is the derivative of \mathcal{A}_ζ at $\zeta = 0$.

The state Φ_t is d -dimensional: $\Phi_t(k)$ is intended to approximate $\mu_t(x^k) - \pi_0(x^k)$ for $1 \leq k \leq d$. The output γ_t is an approximation of $\tilde{y}_t = y_t - \bar{y}^0$.

2.3 Local control: inverse filter design

Fig. 3 shows the nominal behavior of an air-conditioning load along with the associated Markovian model, whose sample paths are piece-wise constant. The sampling rate r was chosen so that the mean sampling time $1/r$ is much smaller than the nominal period of the load.




Figure 3: Temperature evolution of a TCL.

Consider a collection of 2,000 similar units, each consuming 1kW of power when operating, and 50% duty cycle. Hence, without any coordination, the average power consumption is about $\bar{Y}^0 := 1\text{MW}$.

Denote the total power consumption at time t by Y_t . Equivalently,

$$Y_t = N \sum_x \mu_t^N(x) \mathcal{U}(x),$$

and denote the deviation $\tilde{Y}_t = Y_t - \bar{Y}^0$. Using centralized control, the loads could be coordinated so that \tilde{Y}_t tracks a square wave of amplitude 1MW nearly perfectly. If the frequency of this square wave is chosen to be the nominal period (approximately 30 mins in this example), then each load would appear to be evolving without external influence. If the frequency is far from the nominal, then the load will receive poor QoS: either excessive cycling, or poor temperature control. The decentralized control strategy described here is designed to respect these constraints.

In every design considered, it is found that the aggregate dynamics exhibit a resonance near the nominal frequency. Fig. 4 shows a Bode plot for a linearized TCL model with transfer function G_ℓ , with resonance at $f_r = 3 \times 10^{-3}$ rads/sec (consistent with a 30 min period).

These observations are motivation for restricting the bandwidth of service from each load to a neighborhood of this resonance, and introducing pre-filtering to flatten the resonance.




Figure 4: Linearized mean-field dynamics for a TCL with and without inverse-filter

For each load, an associated mean-field model and its linearization can be computed exactly. In each example that we have considered, these dynamics are minimum phase, which simplifies the *inverse filter design* proposed here. The outcome of this design is a prefilter that removes the resonance, and makes the linearized dynamics appear all-pass within a prescribed bandwidth.

Let G_ℓ denote the transfer function for the linearized mean-field model, and let M_ℓ denote the pre-filter. The goal is to design the pre-filter so that $M_\ell(jf)G_\ell(jf) \approx 1$ for a range of $f \in \mathbb{R}$. This goal can be re-cast as the robust control problem described next.




Figure 5: Feedback control system for Inverse Filter Design.

Fig. 5 shows a feedback control loop in which K_ℓ is a transfer function to be designed. In the standard robust control framework, one objective for design is to ensure that the transfer function

Par.	AC	Fast WH	Slow WH
θ^{set}	18–22	48–52	48–52
δ	0.8–1	2.95–3	3.95–4
Θ^a	30–34	19–21	19–21
RC	3.5–4.5	30–36	67–73
ϱ	14/2.5	5/1	5/1

Table 1: TCL parameters for ACs and water heaters.

from ζ_t to y_t is nearly unity in some frequency range. This closed loop transfer function is equal to $L_\ell/(1 + L_\ell)$, where $L_\ell = K_\ell G_\ell$. Consequently, a solution to the robust control problem provides an inverse filter design solution $M_\ell = K_\ell/(1 + L_\ell)$. The desired approximation $M_\ell G_\ell = L_\ell/(1 + L_\ell) \approx 1$ is obtained in the specified frequency band.

The robust control problem is posed as an optimization problem over transfer functions. The optimal transfer function K_ℓ is obtained numerically in Matlab using the `mixsyn` command [1, 12].

2.4 Design with heterogeneous loads

We conclude this section with a few details required to incorporate multiple heterogeneous loads in the demand dispatch model.

First, observe that the inverse design whose linearization is plotted in Fig. 4 may result in a local control algorithm that is too aggressive for a TCL load. Without the inverse filter, the gain from an aggregate of these TCLs falls quickly for frequencies $f > f_r$, which suggests a problem with this inverse filter design: excessive cycling of individual loads will occur if the aggregate tracks high frequencies with significant magnitude. On the other hand, capacity of low-frequency tracking is small because of the temperature constraints associated with TCL hysteresis.

Therefore, it is essential to introduce a second filter to restrict the bandwidth to a range appropriate for the corresponding class of loads. Specifics are provided in the experimental results surveyed in the next section.

This section is concluded with a brief description of the nominal behavior of a TCL, and a summary of the parameters used in this paper when considering a collection of heterogeneous loads.

A common model for temperature evolution is the first order differential equation,

$$\frac{d}{dt}\Theta_t = -\frac{1}{RC}(\Theta_t - \Theta_t^a + \theta^g m_t) + W_t, \quad (8)$$

in which Θ_t is the internal temperature, Θ_t^a is ambient temperature, C is thermal capacitance, R is thermal resistance, and W_t models disturbances.

The temperature gain parameter is $\theta^g = R\rho_{tr}$, where ρ_{tr} is the energy transfer rate: ρ_{tr} is positive for TCLs providing cooling, and negative otherwise. The power consumption is the ratio $\varrho = |\rho_{tr}|/\text{COP}$, where the denominator is known as the *coefficient of performance*. The nominal behavior is defined by a temperature set-point θ^{set} and a dead-band range δ , so that $\Theta_t \in [\theta^{set} - \delta/2, \theta^{set} + \delta/2]$. Temperature is regulated to this band via the binary-valued process $\{m_t\}$, whose behavior is defined by hysteresis, as illustrated in Fig. 3.

Table. 1 displays the ranges of values of TCL parameters for air-conditioners and electric water heaters (a subset of those surveyed in [15]). The value of RC is in units of time (hrs). The last row denotes the maximal power consumption, $\varrho = |\rho_{tr}|/\text{COP}$ (kW).

The experiments that follow are based on a heterogeneous collection of loads in which the parameters for the TCLs take on values within these limits.

3 Simulating the Grid

It is found in prior numerical studies that the mean field model accurately matches the dynamics of an aggregate of loads, provided the total number of loads engaged is on the order of hundreds or more [7, 17]. These prior works focused primarily on residential pool pumps; extensions to TCLs are treated in [7], but without any supporting simulations.

Simulations demonstrating the tracking and disturbance-rejection performance along with a cost analysis of demand dispatch are presented in this section. The impact of daily periodic patterns of response from loads is also investigated.

3.1 Design of a virtual battery

The experiments conducted involved four classes of loads: residential air conditioners (AC); small electric water heaters with faster cycle times (f-WH); large electric water heaters with slower cycle times (s-WH); and residential pool pumps. They are distinguished by their nominal period, and based on this, a bandwidth of service was chosen for the design of each bandpass filter. In each case, a second-order Butterworth filter was adopted — the parameters are summarized in Table. 2.




Figure 6: Individual and aggregate load dynamics.

Twenty different subgroups were obtained for each TCL class, through uniform sampling of the values in Table. 1. Each subgroup contains 2,000 loads, implying a total of 40,000 loads in each TCL class. The Markovian model was obtained via Monte-Carlo based on simulations of (8), following Section IV of [2] and prior work. The experiments include 40,000 homogeneous pools with 12 hour cleaning cycles, with a nominal Markov model identical to that used in [7].

For each homogeneous subgroup, the controlled Markov model $\{S_\zeta : \zeta \in \mathbb{R}\}$ was obtained using the myopic design described in Section 2.2. Based on the resulting design, its linearized dynamics were obtained about $\zeta \equiv 0$. This was the basis of the inverse filter design described in Section 2.3. In addition, as described in Section 2.4, each load locally pre-filters the regulation signal using a bandpass filter.

Denote the respective transfer functions for the linearized mean-field model by, respectively, G_{ac} , G_{fwh} , G_{swh} , G_{pl} , and the respective filters (inverse \times band-pass) by M_{ac} , M_{fwh} , M_{swh} , M_{pl} . The linear model of the aggregate dynamics of all the loads is defined by the sum,

$$H_{load} = M_{ac}G_{ac} + M_{fwh}G_{fwh} + M_{swh}G_{swh} + M_{pl}G_{pl} \quad (9)$$




Figure 7: Open-loop tracking with 40,000 heterogeneous TCLs. The lower plots show the on/off state m for a typical load.

The Bode plot for H_{load} is shown in Fig. 6. The rapid decline in the magnitude plot beyond $f = 10^{-2}$ rads/sec is due to the inherent bandwidth constraints of the loads. Hence, the actuation is augmented with an ideal resource $G_a \equiv 1$. A high pass filter M^{HP} was designed with bandwidth beyond $f = 10^{-2}$ rads/sec, so that the introduction of this resource flattens out the Bode plot. The total response is modeled by the transfer function $H_{\text{total}} = H_{\text{load}} + M^{\text{HP}}G_a$, whose Bode plot is also shown in Fig. 6.

The actuation obtained from G_a might come from batteries, responsive generators, or fast responding loads that provide accurate tracking. The time-scales of ancillary service from these resources are assumed to be in the range of primary control, which is why accurate response is needed in this bandwidth.

Load	Period	BW (cyc/hr)	$\bar{\varrho}_{\text{tot}}$	ϱ_{tot}
AC	20min–1hr.	[1, 1/0.2]	97	224
f-WH	2–4 hrs.	[1/3, 1/0.5]	11	200
s-WH	8–12 hrs.	[1/9, 1]	8.5	200
Pools	24 hrs.	[1/24, 1/3]	20	40

Table 2: Load Dynamics and Power Characteristics: BW of M^{BP} ; max and average power ϱ_{tot} , $\bar{\varrho}_{\text{tot}}$ in MW for 40,000 loads.

Have we constructed a perfect battery? Recall that the nonlinear dynamics have been linearized for the sake of analysis, but the aggregate dynamics remain nonlinear. Moreover, the Bode plot for the linearized dynamics with transfer function H_{total} is not entirely flat in magnitude or phase. These shortcomings are no different than what would be expected for a generator providing balancing service, or a realistic (and imperfect) battery system.

The next results illustrate the accuracy of tracking, and the application of the ensemble of loads for balancing the grid.

3.2 Open loop tracking

The balancing reserves deployed (BRD) from the Bonneville Power Administration (BPA) were used as a reference signal to evaluate open-loop tracking. A single typical windy day, February 19, 2016, was chosen in the open-loop experiments described here. These experiments illustrate the input-output behavior of each collection of TCLs.

For each of the three classes of TCLs, the BRD were passed through a bandpass filter designed based on the frequency characteristics of the class. Fig. 7 shows the *open loop* tracking performance in each case (for the case of AC, the plot shows only six hours during the day). The tracking accuracy

is remarkable for a one-way communication architecture from the grid operator to the loads.

We estimate that the AC trajectory represents only 20% of capacity (the signal could be scaled up by 5 while maintaining reasonable tracking), and the other two plots represent about 50% of capacity. While demand dispatch does increase cycling of TCLs, in these experiments, it was found that cycling was increased by only about 5% from nominal. Without the inclusion of “opt out” control, additional cycling will increase as the magnitude of the reference signal increases [6].

The entire BRD signal can be tracked using a combination of pools and the three classes of heterogeneous TCLs along with the high-frequency ideal resources $M^{\text{HP}}G_a$. Results from experiments in non-ideal settings are described next.

3.3 Closed loop performance

Simulations were performed to evaluate the disturbance rejection performance of the demand dispatch control architecture. The experiments were based on the closed loop system represented in Fig. 1.

The simulations were run in continuous time using Simulink. Full details are contained in the Appendix.

The nonlinear mean field model tracks the aggregate of loads perfectly in all cases considered. In particular, in each of the simulation results shown in Fig. 7, the mean-field model output is nearly indistinguishable from the aggregate stochastic output.

Since it is much faster to simulate the nonlinear deterministic system, we see no reason to conduct a stochastic simulation in these experiments.

The demand dispatch simulation model was based on 1 million ACs, 5 million f-wh, 5 million s-wh, and a large number of pools (this number was taken as a parameter in this study). Each group of loads evolves according to the corresponding mean-field model (6), which is linear in the state and nonlinear in the input.

For a homogeneous group ℓ , there is by design a controlled generator $\{\mathcal{A}_\zeta^\ell\}$, and a linear filter M^ℓ that determine local control. The dynamics of the aggregate of loads in this subgroup evolves as

$$\frac{d}{dt}\mu_t^\ell = \mu_t^\ell \mathcal{A}_{\zeta_t}^\ell, \quad \zeta_t^\ell = M_\ell \zeta_t \quad t \geq 0.$$

Given the specified mix of TCL loads and assuming that the closed-loop system is driven by the BRD signal as the disturbance entering the grid, in order to obtain a flat Bode plot for total actuation as shown in Fig. 6, we would require *4 million pools*! With a PI control architecture, a flat response at low frequencies is not necessary, so experiments were conducted with 1 million residential pools (the approximate number of pools in Florida). The maximum load is thus 1GW, and the average load is 500 MW, so that the pools can at best track signals of ± 500 MW. Tracking was poor when the BRD signal exceeded this range.

Other resources such as commercial water chillers could be added to increase capacity at low frequencies. Instead, in the next set of simulations, the pools were augmented with a single 1 GW generator. This was modeled through the introduction of an additional ideal actuator:

$$H_{\text{total}} = H_{\text{load}} + M^{\text{HP}}G_a + \frac{1}{4}M^{\text{LP}}G_a$$

in which the second-order low-pass filter M^{LP} has unity gain at low frequencies, with the exact pole/zero locations used for the pool loads. The scaling of $1/4$ is introduced so that the response from the ideal low frequency actuators is commensurate with the pools. The resulting Bode plot is no longer flat – its gain below 10^{-4} rads/sec is approximately half of the gain above 10^{-3} rads/sec.




Figure 8: Closed-loop tracking with residential air-conditioners, electric water heaters, pool pumps, and ideal actuators. Actuation from the loads can be interpreted as virtual energy storage.

As a result of the gain variations in the linearization and the nonlinearities caused by capacity constraints, the open-loop tracking will no longer be perfect, especially when the BRD signal takes on large values. While imperfect, the performance is still better than what is received from many generation units (such as Fig. 10 of [9]).

The following set of experiments are based on the closed-loop architecture used in practice today: the BA observes frequency deviations (or some other measure of power mismatch), and varies the balancing reserve/AGC signal in response. Results from these experiments are described next.

The plots on the left hand side of Fig. 8 show the resulting closed loop behavior over 5 days, using BPA BRD data from February 19–23, 2016, as the disturbance D entering the grid (modeled as an additive input disturbance as shown in Fig. 1).

The aggregate response from all actuators, U^A , is approximately the negative of the BRD, so that the frequency deviation is tightly controlled: the disturbance rejection performance is nearly perfect. The grid frequency remained within the range 59.993 to 60.007 Hz over the 5 day period.

The plot on the right hand side of Fig. 8 shows the filtered control signal $U_t^{LP} := M^{LP}U_t$ along with two responses: from the collection of pools, and from the 1 GW generator. The response of the pools nearly matches the response from the ideal generator.




Figure 9: Tracking remains perfect even when the gain from ACs is sinusoidal with a 24 hour period and magnitude range $1 \pm \frac{1}{2}$.

3.4 Time-varying capacity

The time-varying nature of many commercial and residential loads is an issue of concern. For example, the number of air conditioners that are in operation, and hence available for ancillary service, is low during the early morning hours and peaks during the late afternoon — see Fig. 9. of [19].

Experiments were conducted in which the gain of the response of the ACs was amplified/attenuated using a time-varying gain function:

$$g(t) = 1 - 0.5 \sin(f_d t), \quad t \geq 0$$

in which $f_d = 727 \times 10^{-7}$ rads/s corresponds to a 24 hour period. All of the other resources were left the same as the simulation setting of Section 3.3. Fig. 9 shows the results. While the AC actuation does not track its input signal ζ^{ac} , the aggregate actuation from all the resources is almost the exact opposite of the disturbance, just as seen in previous experiments. Disturbance rejection is nearly perfect, and the grid frequency remains within [59.993, 60.007] Hz.

The potential cost of these gain fluctuation is additional actuation from other resources [14].

3.5 Resource availability and cost

Following installation of equipment to enable demand dispatch, the operating cost is essentially zero. Consumers may require incentives to participate (e.g., the monthly credits provided by Florida Power and Light through their OnCall program), but they will also receive some guarantees regarding constraints on QoS and potential costs from additional cycling of equipment.

The benefit of demand dispatch from low frequency services such as residential pools is clear: one million pools serve as a substitute for a 1GW generator. Following the initial investment (usually in \$B), a generator requires fuel, maintenance, and staff. The loads provide accurate regulation service without any of these operating costs.




Figure 10: Cost as a function of capacity from AC loads.

What about high-frequency ancillary services? To investigate the value of the highest frequency services from demand dispatch, we consider a parameterized family of models in which the contribution from air conditioners is varied according to the fraction $\alpha \in [0, 1]$. The remaining $1 - \alpha$ of regulation is obtained from ideal actuation from batteries or other sources. Denote the output of the ideal actuators by $\{U_t^a\}$. The total ideal actuation is defined by the sum:

$$U_t^a = [M^{\text{HP}} G_a + (1 - \alpha) M_{\text{ac}}^{\text{BP}} G_a] U_t$$

Recall $G_a \equiv 1$ in these experiments. The second component is thus $(1 - \alpha) M_{\text{ac}}^{\text{BP}} U_t$, which is intended to replace the lost service from the ACs.




Figure 11: Left: Hypothetical CAISO Net-load over one day in 2020, and its frequency decomposition. Right: “Residual Load” = “Net Load” – “Low Pass” is tracked nearly perfectly. The introduction of demand dispatch alongside generation reduces needed generation capacity by at least 5 GW.

Following [14], the mean-square cost of the closed loop system is defined as,

$$J^2 = \frac{1}{T} \int_0^T |U_t^a|^2 dt \quad (10)$$

This is similar to the “mileage” metric used for ancillary service resources such as batteries. Fig. 10 shows a plot of this cost as a function of α for T corresponding to one day; $\alpha = 1$ corresponds to the simulation setting of Section 3.3. The total cost is reduced by more than 50% when $\alpha = 1$ as compared to $\alpha = 0$.

The cost would be much higher for intermediate values of α if the inverse filter was not used to construct ζ^{ac} [14].

3.6 Ramp services

The plot on the right in Fig. 11 shows a stylized “duck curve” representing the net-load at CAISO anticipated in the near future, based on the assumption that there will be significant solar energy penetration. The plot is based on approximately 10 GW of solar power at peak.

The 15GW ramp observed between 3pm and 6pm is of concern today. It is argued in [14] that the ramp can be smoothed by first scheduling generation to track a low-frequency component of the net-load — denoted “low pass” in the figure. The remaining two zero-energy signals shown can be tracked using a combination of resources — batteries, responsive generators, and demand dispatch.

The mid-pass signal remains substantial – a range of ± 5 GW. This signal could be provided using gas turbine generators, but a total capacity of 10GW would be required. This value can be reduced significantly by applying the same techniques used to address the balancing reserves signal.

Let L_t^r denote the residual load, defined as “Net Load” – “Low Pass”. This is plotted on the right in Fig. 11, where it is seen that it takes on values approaching ± 4 GW. The capacity from loads in the previous set of experiments was insufficient to track this signal. The capacity from TCLs was doubled, so that the simulation was based on 10 million s-WH, 10 million f-WH, and 2 million ACs. It included 1.2 million pools (the approximate number of pools in California), and also ± 2 GW of low frequency regulation that might come from generation or demand dispatch from other loads such as water chillers and water pumping (a significant load in California).

The plots of deviation of power from TCLs shown on the right in Fig. 11 are significant, even though the loads themselves do not deviate from their individual temperature setpoints. The variation in power consumption of s-WH and pools helps to address the “mid pass” signal shown on the left of Fig. 11, whereas the “high pass” component is serviced by the f-WH and AC power consumption. The residual load and aggregate actuation match nearly perfectly.

4 Conclusions

It is exciting to see how heterogeneous loads can coordinate through distributed control to smooth out enormous shocks to the grid. The collection of heterogeneous loads is a multi-GW virtual battery capable of impressive actuation in response to the control signal from the grid operator. In a closed-loop setting, the demand dispatch architecture can perform near-perfect disturbance rejection, tightly controlling the grid frequency. Consequently, demand dispatch offers tremendous potential to provide high-quality ancillary services on timescales spanning from several hours to a few minutes (the time-scale of AGC).

Two issues require further attention. First is the role of the “perfect actuators” supplying regulation at time scales of tens of seconds and faster (the timescale of today’s primary control). Can loads assist with this service as well as bolster synthetic inertia? The analysis in [14] suggests that this could bring risk in terms of stability, but this may depend on other elements of the grid architecture (e.g., the number and size of synchronous generators).

A second, far more significant issue is the time-varying nature of many loads. For example, the nominal load from commercial and residential air-conditioning is roughly periodic over a typical week, and its magnitude changes slowly depending upon the weather. The results summarized in Section 3.4 offer significant hope in terms of system stability. Moreover, it is conjectured that periodicity is a benefit in regions with significant solar energy, since demand is in harmony with supply.

Future work is required to convince the scientific community and the power industry that the overall coupled dynamics will not introduce any additional risk when compared to traditional methods for balancing and frequency regulation. Further large-scale simulation is required along with large-scale demonstration projects.

References

- [1] G. Balas, R. Chiang, A. Packard, and M. Safonov. Robust control toolbox. *For Use with Matlab. User’s Guide, Version, 3*, 2005.
- [2] A. Bušić and S. Meyn. Distributed randomized control for demand dispatch. To appear, IEEE Conference on Decision and Control, March 2016.
- [3] D. Callaway and I. Hiskens. Achieving controllability of electric loads. *Proc. of the IEEE*, 99(1):184–199, January 2011.
- [4] D. S. Callaway. Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. *Energy Conversion and Man.*, 50(5):1389–1400, 2009.
- [5] H. Chavez, R. Baldick, and S. Sharma. Regulation adequacy analysis under high wind penetration scenarios in ERCOT nodal. *IEEE Trans. on Sustainable Energy*, 3(4):743–750, Oct 2012.
- [6] Y. Chen, A. Bušić, and S. Meyn. Individual risk in mean field control with application to automated demand response. In *53rd IEEE Conference on Decision and Control*, pp. 6425–6432, Dec 2014.

- [7] Y. Chen, A. Bušić, and S. Meyn. State estimation for the individual and the population in mean field control with application to demand dispatch. *To appear, IEEE Trans. on Auto. Control*, 2016.
- [8] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent. Aggregate flexibility of thermostatically controlled loads. *IEEE Trans. on Power Systems*, 30(1):189–198, 2015.
- [9] B. J. Kirby. Frequency regulation basics and trends. Report prepared for the US DoE – ORNL/TM-2004/291, OAK RIDGE NAT. LAB., 2004.
- [10] S. Koch, J. L. Mathieu, and D. S. Callaway. Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services. In *Proc. 17th Power Systems Computation Conference*, pp. 1–7, 2011.
- [11] P. Kundur. *Power system stability and control*, volume 7 of *EPRI power system engineering*. McGraw-Hill New York, 1994.
- [12] H. Kwakernaak. Robust control and H_∞ optimization tutorial paper. *Automatica*, 29:255–273, 1993.
- [13] R. Malhame and C.-Y. Chong. Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. *IEEE Trans. Automat. Control*, 30(9):854 – 860, Sep 1985.
- [14] J. Mathias, R. Kaddah, A. Bušić, and S. Meyn. Smart fridge / dumb grid? demand dispatch for the power grid of 2020. In *Proc. 49th Annual Hawaii International Conference on System Sciences (HICSS)*. pp. 2498–2507, Jan 2016.
- [15] J. Mathieu, M. Dyson, D. Callaway, and A. Rosenfeld. Using residential electric loads for fast demand response: The potential resource and revenues, the costs, and policy recommendations. *Proc. of the ACEEE Summer Study on Buildings, Pacific Grove, CA*, 2012.
- [16] J. Mathieu, S. Koch, and D. Callaway. State estimation and control of electric loads to manage real-time energy imbalance. *IEEE Trans. Power Systems*, 28(1):430–440, 2013.
- [17] S. Meyn, P. Barooah, A. Bušić, Y. Chen, and J. Ehren. Ancillary service to the grid using intelligent deferrable loads. *IEEE Trans. Automat. Control*, 60(11):2847–2862, Nov 2015.
- [18] M. Roozbehani, M. A. Dahleh, and S. K. Mitter. Volatility of power grids under real-time pricing. *IEEE Trans. on Power Systems*, 27:1926–1940, 2012.
- [19] R. Smith, K. Meng, Z. Dong, and R. Simpson. Demand response: a strategy to address residential air-conditioning peak load in australia. *J. Modern Power Systems and Clean Energy*, 1(3):223–230, 2013.
- [20] S. Tindemans, P. Djapic, J. Schofield, T. Ustinova, and G. Strbac. Resilience performance of smart distribution networks. Tech. report, Imperial College, 2014.
- [21] C. Ziras, E. Vrettos, and G. Andersson. Primary frequency control with refrigerators under startup dynamics and lockout constraints. In *IEEE Power & Energy Society General Meeting*, pp. 1–5. IEEE, 2015.

A Appendix

A.1 Grid-level transfer functions

Table. 3 provides the transfer functions of the macro grid model G_p and the PI compensator G_c used in the closed-loop experiments of Section 3.

LTI System	Transfer Function
G_p	$10^{-5} \frac{2.488s + 2.057}{s^2 + 0.3827s + 0.1071}$
G_c	$516 \frac{s + 0.5}{s}$

Table 3: Continuous-time transfer functions used in the closed-loop simulations.

A.2 Open loop simulations of TCLs

The open loop simulations in this paper involve three groups of TCLs: ACs, f-wh, and s-wh. Each group consists of 20 sub-groups containing 2,000 similar loads in each subgroup.

The temperature evolution of an individual load is defined by the ODE given in (8). The different subgroups are obtained by uniformly sampling the TCL parameters provided in Table. 1.

For a given TCL, a finite state space for \mathbf{X}^n is obtained through the quantization of the interval $[\theta_{min}, \theta_{max}]$, where $\theta_{min} = \theta_{set} - \delta/2$ and $\theta_{max} = \theta_{set} + \delta/2$. For a given integer d , the interval $[\theta_{min}, \theta_{max}]$ is discretized into $d/2$ values as follows:

$$\mathbf{X}^n = \{\theta_{min} + k\theta_{\Delta} : 0 \leq k \leq d/2 - 1\}$$

where $\theta_{\Delta} = (\theta_{max} - \theta_{min}) / (d/2 - 1)$ represents the temperature increments in the interval $[\theta_{min}, \theta_{max}]$. Furthermore, $\mathbf{X}^u(t) = m_t \in \{0, 1\} \equiv \{\ominus, \oplus\}$.

The matrix Q_0 modeling transitions in \mathbf{X}^n is obtained via Monte Carlo simulations of (8).

Let $\{T_k : (k \geq 1)\}$ denote the jump times in a Poisson process with rate r . Let $\{\Delta_k := T_k - T_{k-1} : k \geq 1\}$ and $T_0 = 0$. It follows that Δ_k is i.i.d., with $E[\Delta_k] = 1/r$. Then, Q_0 is identified via Monte Carlo methods as below.

The bivariate distribution from a given state x to a particular temperature state x'_n is obtained as an empirical average,

$$\pi_Q(x, x'_n) = \frac{1}{N} \sum_{k=0}^{N-1} \mathbb{I}(X(T_k) = x, X^n(T_{k+1}) = x'_n),$$

where $x \in \mathbf{X}, x'_n \in \mathbf{X}^n$. Bayes rule motivates the following definition for the transition matrix:

$$\begin{aligned} Q_0(x, x'_n) &= \mathbf{P}\{X^n(T_{k+1}) = x'_n | X(T_k) = x\} \\ &= \frac{\pi_Q(x, x'_n)}{\sum_{x''_n} \pi_Q(x, x''_n)} \end{aligned}$$

Following the notation established in [17], $p^{\oplus}(x_n)$ and $p^{\ominus}(x_n)$ indicate the probability of a TCL unit switching on and switching off, respectively, at temperature state x_n . The transition matrix

R_0 that models the on/off behavior (i.e. transitions in \mathbf{X}^u) is represented as,

$$\begin{aligned} R_0(x, \oplus) &= \begin{cases} 1 - p^\ominus(x_n) & x = (\oplus, x_n) \\ p^\oplus(x_n) & x = (\ominus, x_n) \end{cases} \\ R_0(x, \ominus) &= \begin{cases} p^\ominus(x_n) & x = (\oplus, x_n) \\ 1 - p^\oplus(x_n) & x = (\ominus, x_n) \end{cases} \end{aligned} \quad (11)$$




Figure 12: Probabilities defining R_0 for an AC.

Fig. 12 shows a particular choice of $\{p^\ominus, p^\oplus\}$ for a given AC, in which $p^\oplus(x_n) = p^\ominus(\theta_{max} + \theta_{min} - x_n)$ for each $x_n \in \mathbf{X}^n$.

The Markov transition matrix S_0 can be obtained from (3), and the rate matrix \mathcal{A}_0 can be generated using (2).

In this paper, the TCL temperature interval is quantized into $d/2 = 20$ temperature states. Consequently, $|\mathbf{X}| = d = 40$.

To model the behavior shown in Fig. 3, we set $r = 1/60 \text{ s}^{-1}$ for ACs. Fast water heaters (f-wh) are on for 15 minutes and off for 2.5 hours, while slow water heaters (s-wh) remain on for 25 minutes and off for 9 hours, approximately. To model this asymmetry in on-off durations, we utilize two rate parameters: r_{on} and r_{off} . We set $r_{\text{on}} = 1/40 \text{ s}^{-1}$ and $r_{\text{off}} = 1/500 \text{ s}^{-1}$ for f-wh, and $r_{\text{on}} = 1/40 \text{ s}^{-1}$ and $r_{\text{off}} = 1/1000 \text{ s}^{-1}$ for s-wh.

The BPA balancing reserves (from February 19, 2016) are used as the control signal U . The open-loop simulations are in discrete-time; the TCLs receive U at 20-second intervals. U is passed through a local load-level second-order Butterworth filter; the passband frequency range of this filter is specified in Table. 2. An additional inverse filter is used, as discussed in Section 2.3, in order to obtain a flat input-output response from the aggregate collection of loads in each sub-group; this yields the signal ζ .

The matrix S_ζ is obtained using the myopic design defined in (4). The mean-field behavior of each group of loads is realized via (6). The linear models are specified in (7). Note that the Markov transition matrix for a given discrete-time setting t can be obtained from the rate matrix as $P_\zeta^t = \exp(t\mathcal{A}_\zeta)$, $t \geq 0$, where $P_\zeta^t(x, x') = \mathbf{P}\{X_t = x' \mid X_0 = x\}$ is the transition probability from state x to x' , with $x, x' \in \mathbf{X}$. The stochastic behavior of the loads specified by the transition matrix, i.e. the transition from the current state x to the next state x' for a given load, can be implemented using a uniformly distributed random number generator $\sim \mathcal{U}[0, 1]$.

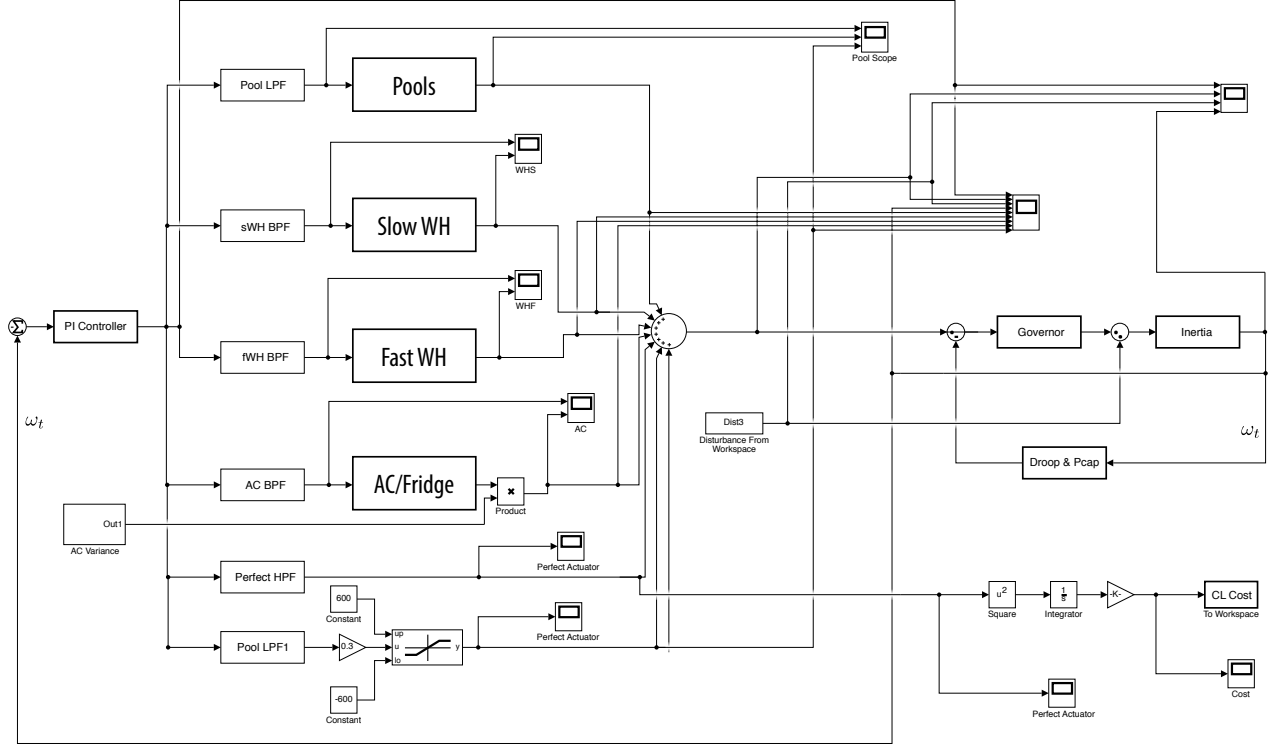


Figure 13: Simulink model.

A.3 Closed loop simulations

The closed loop system shown in Fig. 1 and discussed in Section 3.3 is implemented in continuous time using Simulink. The full Simulink model is illustrated in Fig. 13. Fig. 2 is taken from Fig. 1 of [5], on which this Simulink model is based.

Data from the BPA balancing reserves (February 19–23, 2016) is used for the disturbance D .

The nominal models (S_0, \mathcal{A}_0) of the demand-side resources, which include ACs, f-wh, s-wh, and pools, are generated using open loop simulations (as described in the paper on); the mean-field models are defined by the nonlinear deterministic equation (6), which is non-linear in ζ . The mean-field models are implemented using the Linear Parameter Varying (LPV) block in Simulink, with ζ as the parameter.

The bandpass and inverse filters are implemented as continuous-time linear systems and are similar to the ones used in the open-loop simulations; the bandpass filters are second-order Butterworth filters with frequency ranges provided in Table. 2. The grid-level transfer functions G_p and G_c given in Table. 3 are also implemented using Simulink’s continuous-time LTI system block.